

CLAIMS

1. A method of processing information representing a graph, comprising:
 - serializing each of multiple statements;
 - using a digital processor to independently compute a hash value for each of the multiple statements; and
 - applying a commutative function to each hash value, to obtain an aggregate hash value representing all of the multiple statements.
2. A method according to claim 1, further comprising:
 - identifying at least one attribute of the multiple statements; and
 - using the digital processor to digitally sign a value dependent upon both of the aggregate hash value and the attribute.
3. A method according to claim 2, wherein said method further comprises:
 - identifying as an attribute the number of statements that have been hashed and that are represented by the aggregate hash value; and
 - digitally signing a value dependent upon both of the aggregate hash value and the number of statements represented by the aggregate hash value.
4. A method according to claim 1, further comprising adding a new statement to the multiple statements by:
 - identifying a blank node;
 - assigning a label to the identified blank node; and
 - creating the additional statement in a manner that sets forth association of the blank node with the label.

5. A method according to claim 4, wherein:

said method further comprises using the digital processor to compute a hash value for the additional statement to thereby create an incremental hash value, and combining the incremental hash value with the aggregate hash value using a function which is both commutative and associative; and

computing a hash value includes applying a SHA-1 hash algorithm to obtain the hash value of each statement.

6. A method according to claim 1, further comprising:

identifying a secret key; and

using the digital processor to digitally sign a value dependent upon both of an aggregate hash value and the secret key.

7. A method according to claim 1, further comprising:

identifying a public key; and

using the digital processor to authenticate the graph, including

using the public key to decrypt a verification hash,

comparing the aggregate hash value to the verification hash, and

determining that the graph is authentic in dependence on the comparison.

8. A method according to claim 1, wherein the function is based upon addition within a finite field.

9. A method according to claim 1, wherein:

applying the function to each hash value includes

initiating the aggregate hash value as a constant,

incrementally applying the function to each hashed statement to create therefrom a new aggregate hash value, and

repeating application of the function until all hashed statements and the constant are represented by the aggregate hash value; and
wherein the function is both associative and commutative.

10. A method according to claim 9, wherein the constant is zero.

11. A method according to claim 9, wherein the constant is a secret key, and repeating the application of the function yields an aggregate hash value that is dependent upon both of the secret key and all hashed statements.

12. A method of processing data representing a graph, the graph being stored on machine-readable media, said method comprising:

determining a first hash value representing the graph;

using a digital processor to compute a second hash value that represents an additional statement; and

computing an aggregate hash value that is a function of each of the first hash value and the second hash value, where the function is commutative.

13. A method according to claim 12, wherein:

determining a first hash value includes

retrieving data representing the graph from machine-readable media and also retrieving a verification hash value from machine-readable media,

verifying authenticity of the data representing the graph by locally computing a hash value based upon retrieved data representing the graph, and determining whether the verification hash value matches the locally-computed first hash value, and

if the authenticity of the data is verified, using the matched results as the first hash value and otherwise generating an error; and

adding a new statement to the data by storing the new statement on the machine-readable media, by computing as the second hash value an incremental hash value representing the new statement, and by storing the aggregate hash value on the machine-readable media in association with the graph, where the aggregate hash value represents addition of the first hash value and the second hash value.

14. A method according to claim 12, further comprising:

retrieving the first hash value from remote machine-readable media;

adding the statement by causing the remote machine-readable media to store the additional statement in association with the graph;

computing as the second hash value an incremental hash value representing the additional statement; and

digitally signing the aggregate hash value and causing the remote machine-readable media to store the digitally signed aggregate hash value in association with the data representing the graph.

15. A method according to claim 12, further comprising:

determining an attribute of all data representing the graph, including the additional statement;

generating an aggregate value dependent upon both the aggregate hash and the attribute; and

digitally signing the aggregate value to thereby generate a digital signature.

16. An apparatus, comprising instructions stored on machine readable media, the instructions when executed causing a digital processor to:

compute a hash of each one of multiple statements representing a graph;

apply a commutative function to each of the multiple hashes to thereby yield an aggregate hash representing all of the multiple hashes; and

digitally sign the aggregate hash.

17. An apparatus according to claim 16, wherein the instructions further cause the digital processor to:

compute at least one identifier selected to represent all of the statements represented by the aggregate hash; and

digitally sign the aggregate hash and the at least one identifier.

18. An apparatus according to claim 16, wherein the apparatus permits one to amend resource description framework (RDF) data stored in a web-accessible data store, and to digitally sign that data, said instructions further causing the digital processor to:

accept a new statement from a user;

compute a hash of the new statement;

combine the hash of the new statement with a hash representing an existing RDF graph using the function, where the function is both commutative and associative, to thereby obtain the aggregate hash;

determine an aggregate number of statements representing the RDF graph together with the new statement; and

digitally sign both the aggregate hash and the number.

19. An apparatus according to claim 16, wherein the apparatus permits one to authenticate RDF data, said instructions further causing the digital processor to:

retrieve a set of statements from a remote data store;

retrieve data representing a verification hash from the remote data store; and

authenticating the statements by individually computing a hash of each statement, combining together the hash for each statement to obtain an aggregate hash, comparing the aggregate hash with the verification hash, and determining authenticity based on the comparison.

20. An apparatus according to claim 19, wherein the instructions further cause the digital processor to authenticate the RDF data by further:

retrieving from the remote data store a digitally-signed identifier, computing a locally-derived identifier based on the aggregate set of statements retrieved from the remote data store, and determining whether the digitally-signed identifier and the locally-derived identifier match; and

determining authenticity based on comparison of both the verification hash with the aggregate hash, and the digitally-signed identifier with and locally-derived identifier.

21. In a system that digitally processes descriptive statements about electronic documents for purposes of verifying authenticity of those statements, the system using a data store that stores at least the statements themselves, a digital processor and a digital signature methodology, an improvement comprising:

computing a hash independently for each descriptive statement, to obtain multiple statement-hashes;

applying a function to the multiple statement-hashes to generate an aggregate hash, wherein the aggregate hash is characterized that its value is order-independent relative to the order in which the multiple statement-hashes are processed; and

digitally-signing the aggregate hash.

22. An improvement according to claim 21, wherein:

application of the function is performed by adding the multiple statement-hashes together;

said improvement further comprises identifying the number of statements that are to be signed; and

digitally-signing includes digitally-signing both the aggregate hash and the number of statements represented by the aggregate-hash.

23. An improvement according to claim 21, further comprising:
- processing blank nodes by adding a new statement to the descriptive statements, where the new statement represents a locally-assigned label associating with local blank node processing; and
 - digitally-signing an aggregate hash that represents all of the descriptive statements including the new statement.
24. An improvement according to claim 21, wherein said improvement is used to authenticate descriptive statements retrieved from a remote data store, including by:
- retrieving a verification hash from the remote data store;
 - computing the aggregate hash based upon the descriptive statements retrieved from the remote data store, without first sorting the statements to reorder them;
 - comparing the aggregate hash with the verification hash; and
 - determining authenticity of the descriptive statements retrieved from the remote data store in dependence upon the comparison.
25. An improvement according to claim 24, further comprising:
- retrieving from the remote data store a digitally-signed identifier representing the number of descriptive statements representing an electronic document;
 - locally-computing an identifier of the number of descriptive statements retrieved from the remote data store; and
 - comparing the digitally-signed identifier with the identifier that was locally computed; and
 - determining authenticity of the descriptive statements retrieved from the remote data store in dependence upon both of the comparison of the aggregate hash with the verification hash, and the digitally-signed identifier with the identifier that was locally computed.

26. A data store, comprising machine-readable media and data stored on the machine-readable media, wherein:

the data includes multiple descriptive statements and an aggregate hash of the multiple descriptive statements;

the aggregate hash is computed using a commutative function applied to a hash independently computed for each one of the multiple descriptive statements;

the data also includes at least one attribute representing all of the multiple descriptive hashes; and

the aggregate hash and the at least one attribute are digitally-signed.

27. A data store according to claim 26, wherein:

the descriptive statements are each stored in triples format; and

the at least one attribute includes the number of multiple descriptive statements represented by the aggregate hash.

28. A data store according to claim 26, wherein:

the aggregate hash is computed by adding together hashes independently computed for each one of the multiple descriptive statements and a secret key; and

the secret key is not stored in the data store.

29. A method of doing business by providing verification services for descriptive statements about electronic resources, said method comprising:

hosting a digital signature, said digital signature based upon

an aggregate hash value created by applying a commutative function to multiple independent hash values, each independent hash value based upon a subset of the descriptive statements, and

at least one attribute representing the set of all descriptive statements represented by the aggregate hash value; and

providing the digital signature in response to remote request that calls for verification of descriptive statements in the hash.

30. A method according to claim 29, wherein:

the number of statements in each subset is one; and

the at least one attribute includes the number of descriptive statements represented by the aggregate hash value.

31. A method according to claim 29, further comprising forced labeling of blank nodes, where each blank node is labeled and a descriptive statement identifying each label is included in the set of descriptive statements.

32. A method according to claim 29, further comprising:

hosting the descriptive statements represented by the hash;

hosting among the descriptive statements an identification statement that identifies at least one attribute selected from the set of creator identity, version number, copyright holder, version date, source identity, original creator identity, revision creator identity and owner identity; and

verifying the identification statement upon request.

33. A method of doing business by providing verification services for descriptive statements about electronic resources, said method comprising:

receiving a digital signature and a set of descriptive statements;

validating whether the digital signature conforms to a set of descriptive statements, including

creating an aggregate hash value by applying a commutative function to multiple independent hash values, each independent hash value based upon a subset of the descriptive statements,

decrypting the digital signature and comparing information represented by the signature with the aggregate hash value, to detect whether there is a match condition; and

providing a verification response upon detection of the match condition, and providing a failure response if a match condition is not detected.

34. A method according to claim 33, wherein:

said method further comprises calculating at least one attribute representing the set of all descriptive statements represented by the aggregate hash value; and

decrypting the digital signature and comparing information includes comparing information represented by the signature with the at least one attribute, to detect whether there is a match condition